

Appl. No. 10/082,746
Amdt. dated 09/30/2005
Reply to Office Action of 06/30/2005

REMARKS

Claims 1 - 52 are pending in the present Application. In the above-identified Office Action, the Examiner rejected Claims 1 - 13 under 35 U.S.C. §101 because they are directed to non-statutory subject matter. Claims 1 - 52 were rejected under 35 U.S.C. §102(e) as being anticipated by Fong et al.

In response to the 35 U.S.C. §101 rejection, Claims 1 - 13 have been amended to include the phrase "computer implemented method." By this amendment, Applicants believe that the rejection is no longer warranted and thus request that it be withdrawn.

Further, to overcome the 35 U.S.C. §102(e) rejection of Claims 1, 14, 27, 40, Applicants have amended the claims to further limit the claims. Support for the added limitations can be found on page 15, lines 3 - 11.

For the reasons stated more fully below, Applicants submit that the claims are allowable over the applied reference. Hence, reconsideration, allowance and passage to issue are respectfully requested.

As stated in the SPECIFICATION, presently, resource management software such workload manager (WLM) and partitioning management software do not interact with each other in a way that combines their capabilities to enable the most efficient and flexible solution to logically partitioned systems (LPARs). For example, a critical application running in a partition that has one processor or CPU may consistently be approaching a maximum amount of CPU usage, even though other processors may sit idly, such

AUS920010887US1

Appl. No. 10/082,746
Amdt. dated 09/30/2005
Reply to Office Action of 06/30/2005

as in the case where a CPU is not assigned to a partition. In this instance, WLM would not be able to provide additional CPU resources on its own, but would require the partitioning management software to assign additional CPUs to the partition. Currently, a system administrator has to repartition the computer system whenever this occurs. Having the system administrator handle this task is not as efficient as it would be if the computer system were to do so automatically. The present invention provides such capability to a computer system.

According to one embodiment of the invention, when it is determined that a workload on a resource in a partition exceeds a maximum threshold, a similar resource is automatically re-allocated to the partition. Further, minimum and maximum percentage usages of the resource by processes running in the partition are also automatically varied.

According to another embodiment of the invention, a workload profile is created for each partition. The workload generally has an amount of workload and a time schedule for each workload. Before a workload is to occur and if the workload on the resources originally assigned to the partition is to exceed a maximum threshold, additional resources are automatically re-allocated to the partition.

The invention is set forth in claims of varying scopes of which Claims 1 and 6, reproduced below in their entirety, are illustrative.

1. A computer implemented method of dynamically re-partitioning a partitioned computer system in response to workloads,

AUS920010887US1

Appl. No. 10/082,746
Amdt. dated 09/30/2005
Reply to Office Action of 06/30/2005

each partition of the partitioned computer system having a plurality of resources and a minimum and a maximum percentage usage for each of the resources by each process being executed in each partition, the computer implemented method comprising the steps of:

determining whether a workload on a resource in a partition exceeds a maximum threshold; and

automatically allocating a similar resource to the partition if it is determined that the workload exceeds the maximum threshold, **said automatically allocating step includes the step of automatically varying the minimum and the maximum percentage usage of the resource by each process executing in the partition.** (Emphasis added.)

Applicants submit that Claim 1, as presently drafted, is not anticipated by Fong et al.

Fong et al., teach a flexible dynamic partitioning (FDP) of resources in a cluster computing environment. According to Fong et al., FDP allocates/reallocates resources to partitions. Particularly, partitioning of resources can be initiated by both application and system triggers. Once dynamic partitioning is triggered, FDP allows a partition to invoke a set of resource allocation/reallocation functions associated with its partition. The reallocation function performs a set of resource matchings and determines the necessary resource movement among partitions.

However, Fong et al. do not teach the step of **automatically varying minimum and maximum percentage usages of a resource by processes running in a partition** after a resource has been reallocated to the partition.

AUS920010887US1

Appl. No. 10/082,746
Amdt. dated 09/30/2005
Reply to Office Action of 06/30/2005

Independent Claim 6, reproduced immediately below, was also rejected under 35 U.S.C. §102(e) as being anticipated by Fong et al. Applicants respectfully disagree.

6. A computer implemented method of dynamically re-partitioning a partitioned computer system in response to workloads, each partition of the partitioned computer system having a plurality of resources, the computer implemented method comprising the steps of:

creating a workload profile for each partition, the profile having a workload and a workload time schedule; and

automatically allocating additional resources to a partition before the workload is to occur if the workload on the resources originally assigned to the partition is to exceed a maximum threshold. (Emphasis added.)

As shown in the above-emboldened-italicized limitations, a workload profile that includes a workload time schedule as well as a workload for each partition is used to determine when resources are to be reallocated to the partitions. Thus, unlike the schedule scheme taught by Fong et al. and which the Examiner seemed to have relied on to reject the claim, the workload of the partitions (from the workload profile) controls when the reallocation is to occur.

By contrast, the workload scheme taught by Fong et al., is directed toward a hierarchical scheduling framework. That is, the partitioned system, as taught by Fong et al., is divided into a hierarchical domain whereby the top-level domain contains two or more partitions, the next to the top-level domain being sub-partitions and the

AUS920010887US1

Appl. No. 10/082,746
Amdt. dated 09/30/2005
Reply to Office Action of 06/30/2005

next to the next top-level domain being sub-sub-partitions etc. According to the scheduling scheme, resources are re-allocated to the top-level domain first and then to the next to the top-level domain etc. (see col. 5, lines 6 - 21). Clearly, this scheduling scheme is quite different from the workload time scheduling used in the present invention.

Since the applied reference does not anticipate neither independent Claims 1 nor 6 and since all the claims in the Application contains the limitations in Claims 1 and 6 in some form or fashion. Applicants submit the claims are allowable. Hence, Applicants once more respectfully request reconsideration, allowance and passage to issue of the claims in the application.

Respectfully submitted,

By: 

Volel Emile
Attorney for Applicants
Registration No. 39,969
(512) 306-7968

AUS920010887US1

Page 18 of 18